IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

VLSI TECHNOLOGY LLC, Plaintiff, v. INTEL CORPORATION, Defendant.	Civil Action No.: 1:19-cv-00977-ADA

DEFENDANT INTEL CORPORATION'S OPENING CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

I. INTRODUCTION

In this *Markman* proceeding, the parties' disputes stem from fundamental disagreements about the proper approach to claim construction. For its part, Intel has proposed to construe a term only where the intrinsic record makes clear that the term does not have its plain meaning. For example, Intel has offered constructions (1) to prevent VLSI from attempting to read the claims as covering elements of the prior art that the patents repeatedly criticize and claim to improve upon, and (2) to hold VLSI to the representations about claim scope that the patentee made to the Patent Office to secure issuance of the claims. And the remaining terms require construction because they are written in means-plus-function format—several of which are indefinite because they require an impossible function or specify a function lacking any corresponding structure.

By contrast, VLSI's proposed constructions largely consist of blanket assertions that "plain meaning" applies, without ever specifying what that meaning might be—presumably so it can later argue that the term means whatever VLSI wants it to mean, regardless of what was said about the term's meaning in the patent and file history. The Federal Circuit has repeatedly criticized this type of tactical approach to claim construction, and it should be rejected here as well.

Accordingly, Intel's proposed constructions should be adopted.

II. U.S. PATENT NO. 6,366,522

A. Background

The '522 patent claims to improve power consumption in an integrated circuit. One aspect of the alleged invention involves a "voltage regulator" that receives an input voltage at one level (e.g., from a power source, such as a battery), and uses regulating circuitry (e.g., switches) to generate an output voltage at a different level. DX-1 ['522 patent], 2:13-16. A voltage regulator that outputs a voltage ***higher*** than the input voltage is a "boost regulator" or "boost converter," and a voltage regulator that outputs a voltage ***lower*** than the input voltage is a "buck regulator,"

"buck converter," or "down converter." DX-18 [Erickson] at 53-55; DX-6 [Apsel] ¶¶ 17-18.

"Inductance" refers to the property of an electric circuit in which a voltage is induced by changing a magnetic field, and an "inductor" is the component commonly used to do so. Boost and buck regulators both use an "inductance," but in different ways and for different purposes. As shown below, a boost regulator uses an inductance ***as part of its input***—i.e., (1) the input voltage received from a power supply (in green) is combined with the voltage induced by the inductor (in red), and (2) the regulating circuitry (in blue) then uses that combined voltage to generate an output voltage higher than the initial power supply voltage:

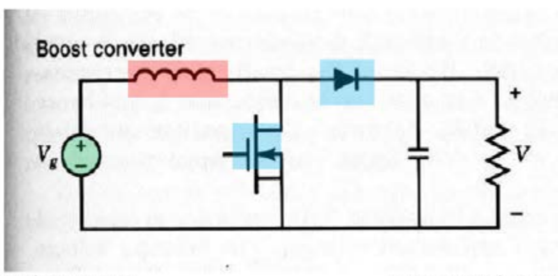

Input voltage: 2 V Output voltage: 3 V

DX-18 [Erickson] at 55, Fig. 4 (annotations added); DX-6 [Apsel] ¶¶ 19, 22-24. By contrast, a buck regulator uses an inductance ***as part of its output***—i.e., (1) the regulating circuitry (in blue) receives an input voltage from a power supply (in green) and generates a lower output voltage, and (2) the inductance (in red) smooths the regulated lower output voltage to a fixed level:

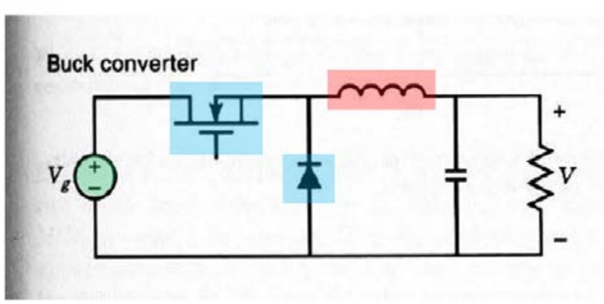

Input voltage: 2 V Output voltage: 1 V

DX-18 [Erickson] at 55, Fig. 4 (annotations added); DX-6 [Apsel] ¶¶ 20-21.

B. Disputed Term: "regulate [regulating] at least one supply from a power source and an inductance" (all claims)

Intel's Construction	VLSI's Construction
regulate [regulating] at least one supply from an inductance connected to a power source, where the inductance is positioned between the power source and the regulating circuitry	plain and ordinary meaning

Intel's construction reflects the applicant's clear and repeated representations to the Patent Office—made to overcome prior art rejections—that this limitation requires a boost regulator configuration in which the claimed "inductance" is connected to the power source and sits between the power source and regulating circuitry. By contrast, VLSI seeks to have the claims cover the very type of buck regulator configuration that the patentee disclaimed to obtain issuance.

Specifically, during reexamination, the Examiner rejected claims in view of eight prior art references. DX-10 [Rejection] at 55-58. As shown below, each reference disclosed a ***buck regulator*** in which the regulating circuity (in blue) receives an input voltage from a power supply (in green), and an inductor (in red) smooths the regulated output received from the regulating circuitry:

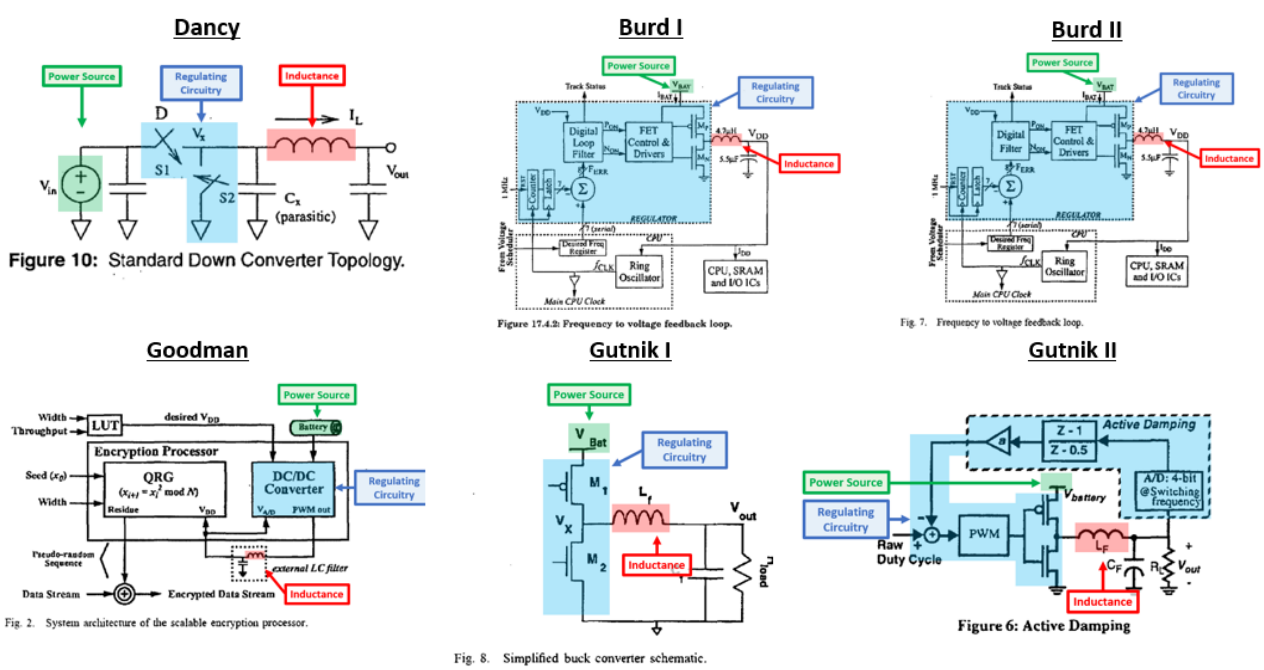

Wei

FIGURE 1. Buck Converter

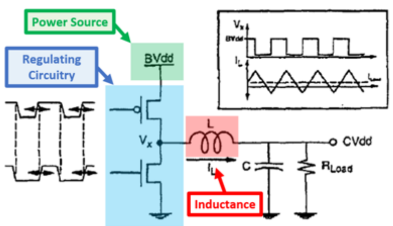

Min

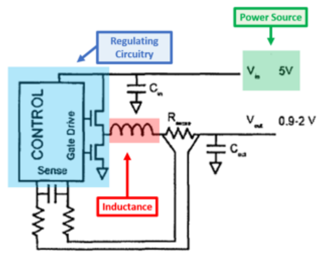

Figure 3. Simplified schematic for the buck regulator.

DX-11 [7/16/2008 Reply] at 16-43 (annotations added); DX-6 [Apsel] ¶ 32.

In response, the patentee argued that "regulating from a power source and an inductance" requires more than a power source and inductance appearing ***anywhere*** in the regulator. Instead, the patentee repeatedly represented that the cited prior art was distinguishable because, consistent with the boost regulator shown in Figure 1 (below), the '522 claims require an inductance ***connected to*** a power source and ***positioned between*** the power source and regulating circuitry:

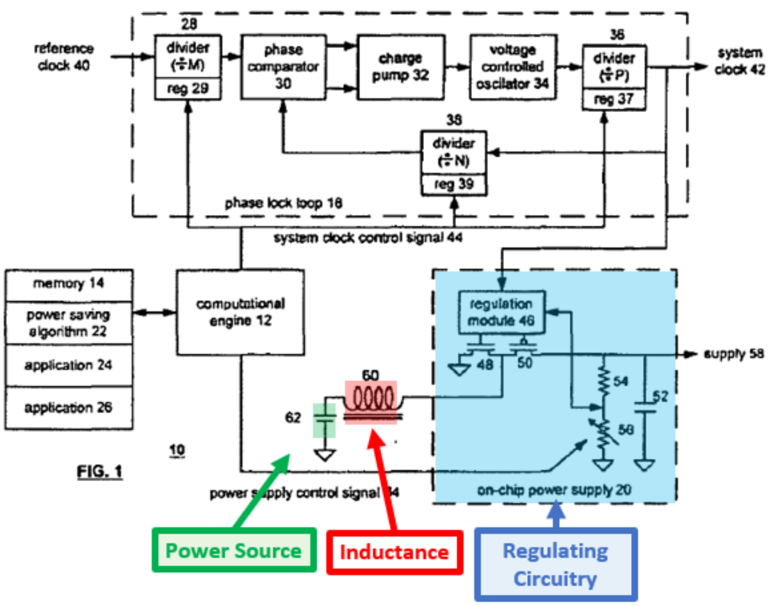

DX-1 ['522 patent], Fig. 1 (annotations added); DX-11 [Reply] at 16 (arguing Figure 1 shows regulating "supply 58 from power source 62 (e.g. a battery) ***and*** inductance 60"); DX-6 [Apsel] ¶ 33.

For example, the buck regulator disclosed in the cited Wei prior art reference (below)

included a power source (in green), regulating circuitry (in blue), and an inductance (in red):

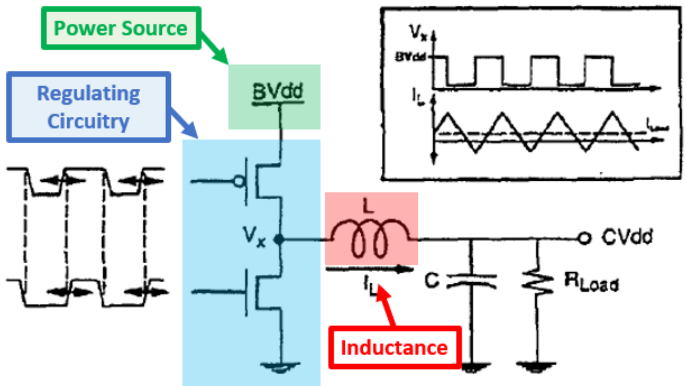

DX-11 [Reply] at 26 (annotations added). To overcome the Examiner's rejection, the patentee represented that this configuration failed to meet the "regulate [regulating] at least one supply from a power source and an inductance" limitation because the inductance was not connected to the power source, as in a boost regulator, and instead appeared ***after*** the regulating circuitry: "Wei's inductance is ***not connected to a power source*** and Wei does not disclose regulating CVdd 'from a power source and an inductance.'" *Id.*; DX-6 [Apsel] ¶¶ 38-39.

Likewise, the cited Min prior art reference disclosed the buck regulator below, in which the inductance (in red) is not connected to the power source (in green), and instead is positioned ***after*** the regulating circuitry (in blue):

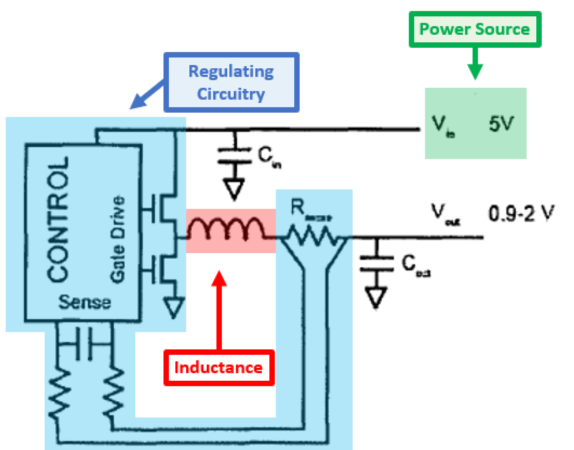

Figure 3. Simplified schematic for the buck regulator.

DX-11 [Reply] at 22-23 (annotations added). The patentee again represented that the claims did not cover a buck regulator configuration because "Min's inductance ***is not connected to a power source***." *Id.* at 22-23; DX-6 [Apsel] ¶¶ 36-37.

Similarly, the cited Dancy prior art reference disclosed regulating circuitry (switches S1 and S2, in blue) that received an input from a power supply (in green) and used an inductor (in red) appearing ***after*** the regulating circuitry for smoothing purposes:

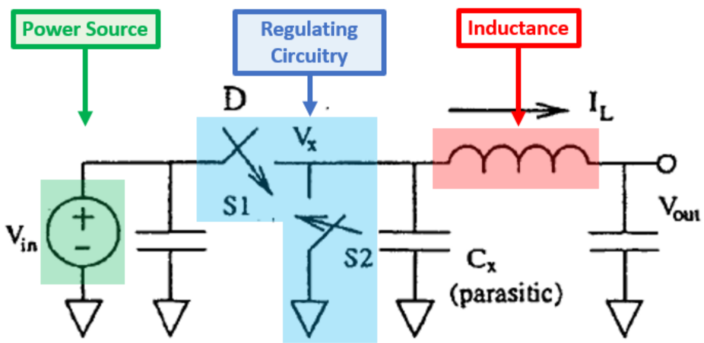

Figure 10: Standard Down Converter Topology.

DX-11 [Reply] at 16 (annotations added). The patentee argued that "the power supply control signal ***operates switches S1 and S2 from the power source (Vin) alone***, not from the power source ***and*** the inductance as recited in claim 1." *Id.* Thus, the patentee again made clear that the "inductance" required by the '522 claims must sit between the power source and regulating circuitry—and not after the regulating circuity, as in a buck regulator. DX-6 [Apsel] ¶¶ 34-35.

The patentee similarly distinguished the five other cited prior art references because the inductances in the disclosed buck regulators appeared ***after*** the regulating circuitry, and thus were ***not connected*** to the power source and ***did not sit between*** the power source and regulating circuitry. DX-11 [Reply] at 29 (distinguishing Goodman's "buck configuration" because the inductance in the "'external LC filter' is ***not connected to a power source***"); *id.* at 32, 36 (distinguishing Burd I and II because the "external LC filter" is in a "buck configuration"); *id.* at 39 (distinguishing Gutnik I because the inductance is in "a buck converter" configuration); *id.* at

42-43 (distinguishing Gutnik II because it "uses the LC filter in a buck configuration"); DX-6 [Apsel] ¶¶ 40-49.

These many representations make clear that the '522 claims require a boost regulator configuration in which the "inductance" is connected to the power source and sits between the power source and the regulating circuitry. DX-6 [Apsel] ¶¶ 50-52. Because VLSI is bound by those representations, Intel's construction should be adopted. *See Comput. Docking Station Corp. v. Dell, Inc.*, 519 F.3d 1366, 1376 (Fed. Cir. 2008) (construing "portable computer" to mean "a computer without a built-in display or keyboard that is capable of being moved or carried about" because, during prosecution, "the applicants clearly distinguished their invention from computers with a built-in display or keyboard"); *Krippelz v. Ford Motor Co.*, 667 F.3d 1261, 1266 (Fed. Cir. 2012) (construing claim term to include requirements argued by patentee during reexamination).

III. U.S. PATENT NO. 7,292,485

A. Background

The '485 patent relates to writing data to memory cells. As the patent's Abstract and background section explain, a "memory array" typically includes (1) one or more "lines" of memory cells, in which each line has multiple memory cells and each cell holds one bit of data, and (2) conductors and other circuitry that supply voltage to power the memory cells. DX-2 ['485 patent], Abstract, 1:12-17; DX-6 [Apsel] ¶ 54. As was known in the prior art, supplying a higher voltage to a memory cell improves the cell's stability, making data loss less likely. DX-2 ['485 patent], 1:31-35. However, a higher voltage also makes it more difficult to store (i.e., "write") new data to the memory cell. *Id.*, 1:36-41. One known way to resolve these conflicting considerations was to ***lower*** the voltage supplied to memory cells during write operations to those cells, a technique known as a "write assist." DX-6 [Apsel] ¶¶ 53-55.

The '485 patent describes a specific write assist technique in which, when writing to a line

of memory cells, the voltage to that memory cell line is lowered by sharing charge between the memory cells and a line of "dummy cells." For example, Figure 2 below depicts a memory array with a first line of memory cells 13 (in green), a second line of memory cells 15 (in blue), and a "capacitance structure" including a line of dummy cells 17 (in red):

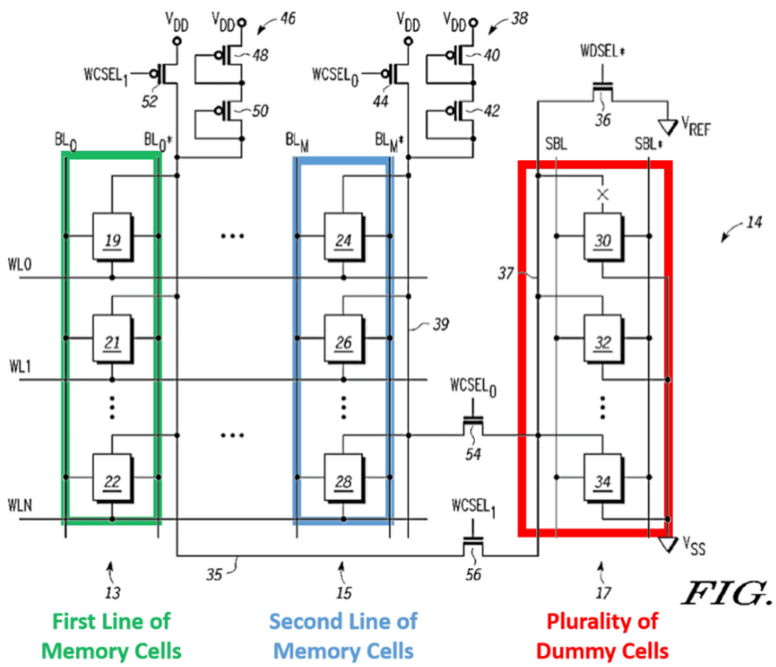

FIG. 2

DX-2 ['485 patent], 3:10-15, 3:59-62, Fig. 2 (annotations added). According to the patent, dummy cells are like "conventional" memory cells—but instead of storing data, the dummy cells are used for charge sharing. *Id.*, 3:61-63, 4:39-46, 5:41-43, 6:22-32.

The patent describes (and claims) embodiments involving a write operation to the second line of memory cells. During that operation, the first line of memory cells (to which data is not being written) remains coupled to its power supply (terminal V_{DD}) through its power supply line 35 (in solid yellow), and thus continues to receive the higher voltage. During the write operation, however, the second line of memory cells is ***decoupled*** from the power supply voltage (shown by red "X" below) and ***coupled*** instead to the line of dummy cells (through transistor 54, in orange):

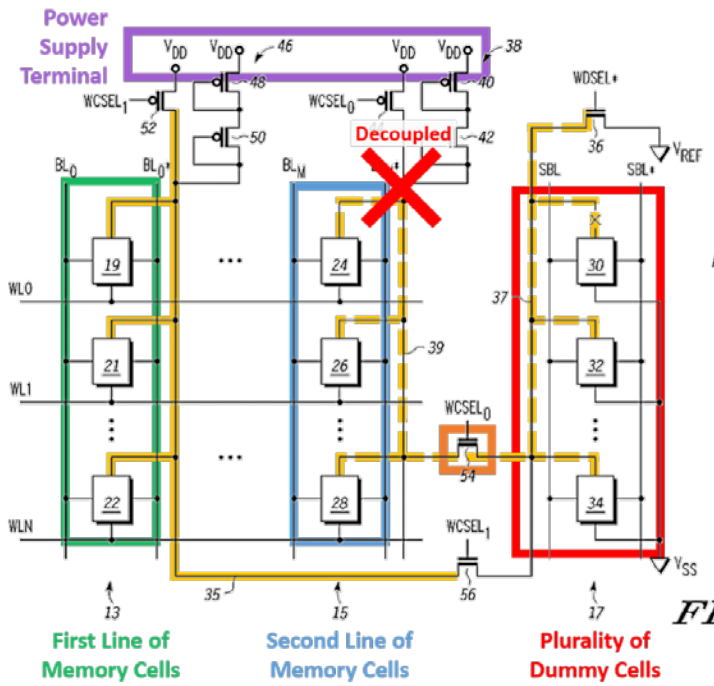

FIG. 2

DX-2 ['485 patent], 4:33-44; 6:14-27, Fig. 2 (annotations added). This decoupling (from the power supply voltage) and coupling (to the dummy cells) causes electrical charge to be shared between the second line of memory cells and the dummy cells through their conductor lines 37 and 39 (in dashed yellow), thereby reducing the voltage supplied to the second line of memory cells during the write operation to those cells. The amount by which the voltage is reduced depends on the relative "capacitance" (the ability to store electrical charge) of the second line of memory cells and the line of dummy cells. DX-6 [Apsel] ¶ 58.

B. Disputed Term: "a capacitance structure" (claims 1, 12, 17)

Intel's Construction	VLSI's Construction
Function: providing "capacitance" **Structure**: (1) Dummy column 17 (comprised of dummy SRAM cells 30, 32, and 34 and dummy bitlines SBL and SBL*) and conductor 37, configured to be selectively coupled to one or more of the dummy SRAM cells, as shown in Figure 2, and equivalents thereof; or alternatively (2) dummy row 70 (comprised of dummy SRAM cells 82, 84, and 86 and dummy wordline SWL) and conductor 71, configured to be selectively coupled to one or more dummy of the SRAM	Plain meaning; not means-plus-function

cells, as shown in Figure 3, and equivalents thereof	

The parties dispute whether "capacitance structure" should be (1) construed as a mean-plus-function term, as Intel contends, or (2) given no construction, as VLSI contends. Because the term is purely functional and does not connote any definite structure, Intel's construction applies.

First, 35 U.S.C. § 112, ¶ 6 applies where a claim element is "expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof." The Federal Circuit has made clear that, "[w]hen a claim term lacks the word 'means,' … § 112, para. 6 will apply if … the claim term fails to 'recite sufficiently definite structure' or else recites 'function without reciting sufficient structure for performing that function.'" *Williamson v. Citrix Online, LLC*, 792 F.3d 1339, 1349 (Fed. Cir. 2015) (*en banc*). For that reason, Section 112, ¶ 6 applies to generic "nonce" words—like "element," "device," "mechanism," or "module"—that are "tantamount to using the word 'means.'" *Id.* at 1350.

That is precisely the case here. The term "capacitance structure" does not connote any specific structure to persons of ordinary skill in the art. Instead, the term merely recites a "structure" for performing the function of providing "capacitance." DX-6 [Apsel] ¶¶ 65-66. The generic, nonce word "structure" connotes no specific structure, and thus "sets forth the same black box recitation of structure … as if the term 'means' had been used." *Williamson*, 792 F.3d at 1350. Nor does the prefix "capacitance" impart any definite structure. It indicates only that the claimed "structure" functions to provide capacitance; but an ordinarily-skilled artisan would appreciate that virtually ***every*** integrated circuit component has some amount of capacitance (parasitic or otherwise). DX-6 [Apsel] ¶¶ 65-66.

The term "capacitance structure" accordingly is no more than a "generic description" of a component defined in purely functional terms, and thus should be treated as a means-plus-function term. *See Williamson*, 792 F.3d at 1350–51 ("distributed learning control module" is means-plus-

function limitation because it "do[es] not describe a sufficiently definite structure"); *Scarborough v. Integricert, LLC*, 2015 WL 5099128, at *13 (W.D. La. Aug. 31, 2015) ("attachment structure" is mean-plus function limitation because "the phrase fails to recite definite structure"); *In re Katz Interactive Call Processing Patent Litig.*, 2008 WL 4865032, at *13 (C.D. Cal. Mar. 4, 2008) ("gang holding structure" is means-plus-function limitation).[1]

VLSI's position that "capacitance structure" should be given its "plain meaning" ignores that the term has ***no*** understood meaning for any particular structure(s)—as confirmed by VLSI's failure to identify any such meaning. VLSI's position also would improperly render the term "capacitance" superfluous—because, as noted above, virtually ***every*** integrated circuit component has some capacitance, and thus ***any*** circuit structure could be a "capacitance structure." *See, e.g.*, *Akzo Nobel Coatings, Inc. v. Dow Chem. Co.*, 811 F.3d 1334, 1340 (Fed. Cir. 2016) (rejecting construction of "pressurized collection vessel" under which any pressurized vessel would be a "pressurized collection vessel"). DX-6 [Apsel] ¶ 65.

Second, where Section 112, ¶ 6 applies, the scope of a means-plus-function term is limited to "the corresponding structure, material, or acts described in the specification and equivalents thereof." 35 U.S.C. § 112, ¶ 6. Here, Intel's proposed construction identifies the ***only structures*** that the '485 specification discloses as corresponding to the claimed function of providing

[1] *See also Grecia v. Samsung Elecs. Am., Inc.*, 2019 WL 3940193, at *2-4 (Fed. Cir. Aug. 20, 2019) ("customization module" is means-plus-function term because "module" is a nonce word and "customization" "at best describes the module's intended functionality"); *Diebold Nixdorf, Inc. v. ITC*, 899 F.3d 1291, 1298 (Fed. Cir. 2018) ("cheque standby unit" is means-plus-function term because the claims recited no structure and described the term "solely in relation to its function and location in the apparatus"); *Advanced Ground Info. Sys., Inc. v. Life360, Inc.*, 830 F.3d 1341, 1347–48 (Fed. Cir. 2016) ("symbol generator" is means-plus-function term, even though individual words had a known meaning, because "the combination of the terms … suggests that it is simply an abstraction that describes the function being performed"); *Alarm.com, Inc. v. SecureNet Techs., LLC*, 2019 WL 3996883, at *4–6 (D. Del. Aug. 23, 2019) ("connection management component" is means-plus-function term because "component" is a nonce word).

capacitance—two embodiments, each using dummy cells to provide capacitance for charge sharing when writing to a selected line of memory cells. DX-6 [Apsel] ¶¶ 67-68.

The first structure (shown in Figure 2) includes dummy column 17 "used for capacitance sharing" with selected columns in a memory array. DX-2 ['485 patent], 4:3-7, 4:39-55; *id.*, 3:59-65 ("Dummy column 17 … includes a pair of dummy bit lines labeled 'SBL' and 'SBL*.' Dummy SRAM cells 30, 32 and 34 are coupled to the dummy bit lines SBL and SBL* and are conventional SRAM cells …. Each of the dummy cells has a supply terminal that can be coupled to a conductor 37."); DX-6 [Apsel] ¶ 67. And the second structure (shown in Figure 3) discloses an alternative dummy row 70 for "charge sharing" with selected rows of a memory array. DX-2 ['485 patent], 6:22-31; *id.*, 5:39-45 ("Dummy row 70 … includes a word line labeled 'SWL' and all of the cells coupled to SWL. Dummy SRAM cells 82, 84, and 86 are coupled to the dummy word line SWL and are conventional SRAM cells …. Each of the dummy cells has a supply terminal that can be coupled to a conductor 71."); DX-6 [Apsel] ¶ 68.

Intel's construction should be adopted because it captures the structures disclosed in these passages—which are the ***only*** structures the patent discloses as performing the claimed function of providing capacitance. *See Williamson*, 792 F.3d at 1352 (construction of means-plus-function term must identify "corresponding structure" that "the intrinsic evidence clearly links or associates … to the function recited in the claim"); *Medtronic, Inc. v. Advanced Cardiovascular Sys., Inc.*, 248 F.3d 1303, 1315 (Fed. Cir. 2001) (limiting corresponding structure to the structure the patent clearly links to the claimed function, even though a skilled artisan would understand that other non-disclosed structure also could perform the claimed function).

C. Disputed Term: "precharging means for precharging the capacitance structure to a predetermined voltage prior to a write operation for the second line of memory cells" (claim 17)

Intel's Construction	VLSI's Construction
Function: "precharging the capacitance structure to a predetermined voltage prior to a write operation for the second line of memory cells" **Structure**: (1) voltage source V_{REF} and transistor 36, coupled in series to provide a reference voltage to one or more dummy cells through conductor 37, as shown in Figure 2, and equivalents thereof; or alternatively (2) voltage source V_{REF} and transistor 90, coupled in series to provide a reference voltage to one or more dummy cells through conductor 71, as shown in Figure 3, and equivalents thereof	**Function**: "precharging the capacitance structure to a predetermined voltage prior to a write operation for the second line of memory cells" **Structure**: "a conductor, or equivalents thereof"

The parties agree that the "precharging means" is a mean-plus-function term, but dispute the corresponding structure. Intel's proposed structure should be adopted.

As noted above, a means-plus-function term is limited to the structure that the specification "clearly links or associates" with the claimed function. *Williamson*, 792 F.3d at 1352. Here, Intel has identified the precise structure that the specification identifies as performing the claimed function of "precharging the cells of dummy column 17" (or, alternatively, "precharging the cells of dummy row 70"). DX-2 ['485 patent], 4:30-33, 6:13-16. As shown in Figure 2, this precharging occurs when a reference voltage (V_{REF}) is provided to dummy cells 30, 32, 34 (through conductor 37 and transistor 36):

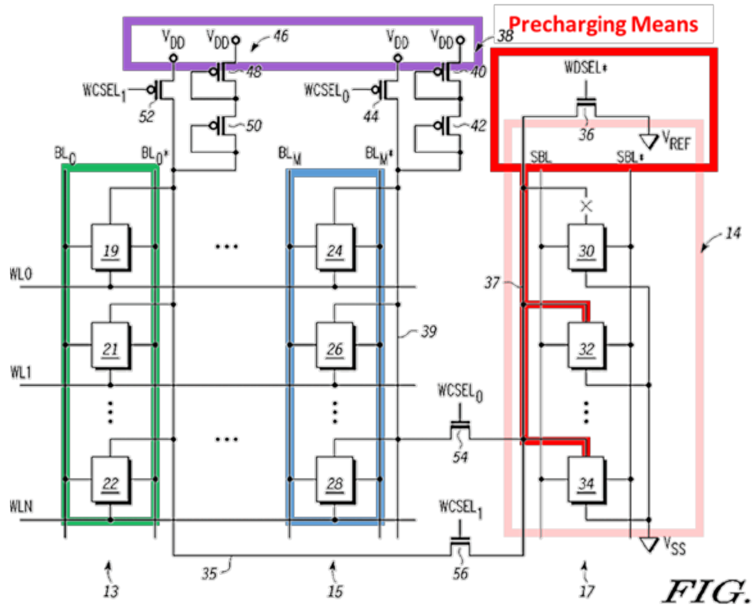

FIG. 2

DX-2 ['485 patent], Fig. 2 (annotations added); *id.*, 4:7-14 ("The conductor 37 is coupled to receive a reference voltage labeled 'V_{REF}' via an N-channel transistor 36."); *id.*, 4:28-33 ("[P]rior to a write operation, control signal WDSEL* is provided as a logic high voltage to cause transistor 36 to be conductive. Reference voltage V_{REF} is provided to the supply terminals of each of the memory cells 30, 32, and 34 precharging the cells of dummy column 17 to V_{REF} (ground)."); *id.*, 5:61-64, 6:11-16 (comparable structure for dummy row 70 in Fig. 3); DX-6 [Apsel] ¶¶ 72-73.

VLSI's proposed structure contains a vague reference to "a conductor." But because a "conductor" is ***any*** component through which charge can flow, that construction would effectively render the term meaningless—and allow it to be met by any of the multiple conductors referenced in the specification with ***no role*** in precharging dummy cells. DX-6 [Apsel] ¶¶ 75-77. *See B. Braun Med., Inc. v. Abbott Labs.*, 124 F.3d 1419, 1425 (Fed. Cir. 1997) ("valve seat" not required structure given the "lack of association between the valve seat and the recited function").

Moreover, even if VLSI were to limit its proposed structure to conductors 37 and 71, those conductors could not, by themselves, perform the claimed "precharging" function. A conductor—without other circuitry—cannot precharge anything; it merely links a voltage source that provides

the voltage needed for precharging and the dummy cells that are precharged. DX-6 [Apsel] ¶ 76. For example, in Figure 2, conductor 37 links reference voltage V_{REF} and dummy cells 30, 32, 34; but precharging occurs only when a control signal is provided "to cause transistor 36 to be conductive" so that the reference voltage is supplied to the dummy cells. DX-2 ['485 patent], 4:5-8, 4:28-33. Thus, to "precharge" the dummy cells, the structure must include conductor 37 as well as reference voltage V_{REF} (which provides the reference voltage) and N-channel transistor 36 (which couples V_{REF} to conductor 37). DX-6 [Apsel] ¶¶ 71-73.

D. Disputed Term: "first coupling means for coupling the power supply terminal to the first power supply line during the write operation for the second line of memory cells" / "second coupling means for coupling the second supply line to the first capacitance structure during the write operation for the second line of memory cells" (claim 17)

'485 Patent Term	Intel's Construction	VLSI's Construction
"first coupling means for coupling the power supply terminal to the first power supply line during the write operation for the second line of memory cells" (claim 17)	**Function**: "coupling the power supply terminal to the first power supply line during the write operation for the second line of memory cells" **Structure**: (1) transistor 52 and clamping circuit 46, configured to couple power supply voltage V_{DD} and conductor 35, as shown in Figure 2 (if the second line of memory cells is SRAM column 15), and equivalents thereof ; or alternatively (2) transistor 96 and a clamping circuit, configured to couple power supply voltage V_{DD} and conductor 67, as shown in Figure 3 (if the second line of memory cells is SRAM row 68), and equivalents thereof	**Function**: "coupling the power supply terminal to the first power supply line during the write operation for the second line of memory cells" **Structure**: "a switching circuit, or equivalents thereof"
"second coupling means for coupling the second supply line to the first capacitance structure during the write operation for the second line of memory cells"	**Function**: "coupling the second supply line to the first capacitance structure during the write operation for the second line of memory cells" **Structure**: (1) Transistor 54, configured to couple conductor 39 to conductor 37, as shown in Figure 2 (if the second line of memory cells is SRAM column 15); or alternatively (2) transistor 94, configured to couple conductor 69 to	**Function**: "coupling the second supply line to the first capacitance structure during the write operation for the second line of memory cells" **Structure**: "a

(claim 17)	conductor 71, as shown in Figure 3 (if the second line of memory cells is SRAM row 68)	switching circuit, or equivalents thereof"

The parties agree that "first coupling means" and "second coupling means" are means-plus-function terms, but dispute the corresponding structures. Intel's proposals should be adopted.

For both terms, Intel's construction identifies the structures that the specification discloses for performing the claimed "coupling" functions, as shown in Figure 2 below:

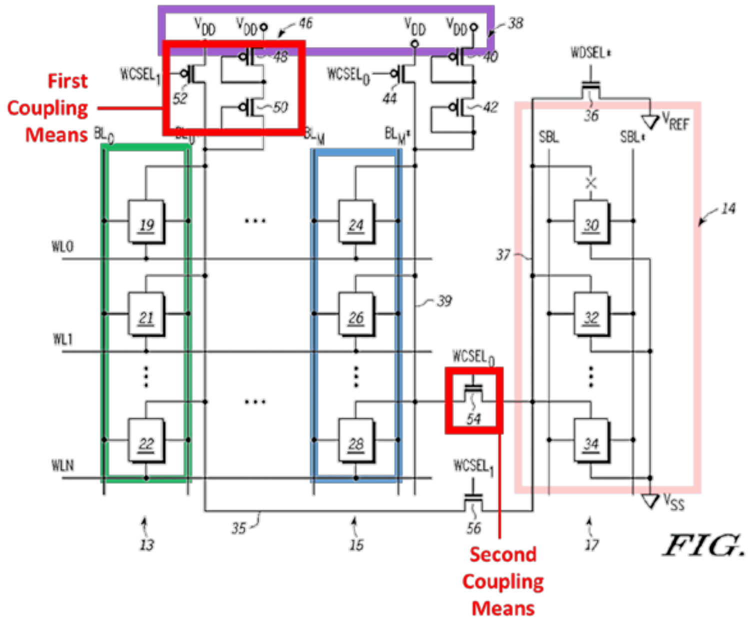

FIG. 2

DX-2 ['485 patent], Fig. 2 (annotations added).

For the "first coupling means"—which must "coupl[e] the power supply terminal to the first power supply line during the write operation for the second line of memory cells"—the specification discloses a transistor and clamping circuit that couple the power supply voltage terminal to the conductor supplying the first line of memory cells and limit any voltage drop on the conductor to a predetermined level. DX-2 ['485 patent], 3:45-52 ("A power supply voltage is selectively provided to the conductors 35 and 39. P-channel transistor 52 has a source coupled to a power supply voltage terminal labeled 'V_{DD}', a gate for receiving a control signal labeled '$WCSEL_1$', and a drain coupled to conductor 35. Clamping circuit 46 includes diode-connected

P-channel transistors 48 and 50 coupled in series between V_{DD} and conductor 35."); *id.*, 5:22-30 (comparable structure for Figure 3), 5:35-40; DX-6 [Apsel] ¶¶ 80-83.

For the "second coupling means"—which must "coupl[e] the second supply line to the first capacitance structure during the write operation for the second line of memory cells"—the patent discloses a transistor that couples the conductor supplying the second line of memory cells to the conductor supplying the dummy cells. DX-2 ['485 patent], 4:15-23 ("In FIG. 2, transistor 54 is used to couple conductor 37 to conductor 39 of column 15 in response to a control signal labeled '$WCSEL_0$' …."); *id.*, 4:33-39, 5:65-6:6, 6:16-22 (comparable structure for Figure 3); DX-6 [Apsel] ¶¶ 86-89.

In contrast, VLSI's proposed structure ("a switching circuit") is not shown in ***any*** figure or described in ***any*** embodiment. Instead, the only reference to a "switching circuit" simply repeats the functional language of claim 1. DX-2 ['485 patent], 6:53-55. Because nothing in the patent describes a "switching circuit" or "clearly links" it to the claimed function, it would be error to include such a component as corresponding structure. *Williamson*, 792 F.3d at 1352.

E. Disputed Term: "decoupling means for decoupling the first power supply line from the second line of memory cells during the write operation for the second line of memory cells" (claim 17)

Intel's Proposed Construction	VLSI's Construction
Function: "decoupling the first power supply line from the second line of memory cells during the write operation for the second line of memory cells" **Structure**: None, indefinite	**Function**: "decoupling the first power supply line from the second line of memory cells during the write operation for the second line of memory cells" **Structure**: "a switching circuit, or equivalents thereof"

The parties agree that the "decoupling means" is a means-plus-function term, but dispute whether the specification discloses any corresponding structure. Because the patent discloses ***no*** corresponding structure for performing the claimed function of "decoupling the first power supply line from the second line of memory cells," the term is indefinite.

The specification never describes an example in which the first power supply line is ever ***coupled*** to the second line of memory cells; thus, it also ***never*** describes ***decoupling*** the first power supply line from the second line of memory cells (much less structure for doing so). DX-6 [Apsel] ¶ 92. Instead, as shown in Figure 2 below, in every example, only the ***first line of memory cells*** (in green) is coupled to the ***first power supply line*** 35 (in yellow); while the ***second line of memory cells*** (in blue) is coupled to its own, separate power supply line—the ***second power supply line*** 39 (not highlighted):

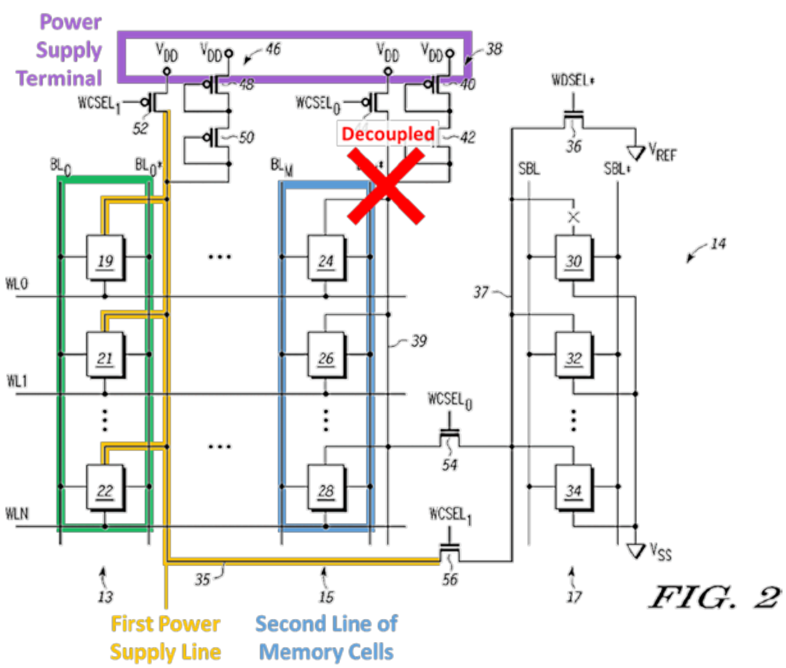

DX-2 ['485 patent], Fig. 2 (annotations added). No disclosed structure is used to "decoupl[e]" the ***first power supply line*** from the ***second line of memory cells***. DX-6 [Apsel] ¶ 92. The term therefore is indefinite. *See Ergo Licensing, LLC v. CareFusion 303, Inc.*, 673 F.3d 1361, 1363 (Fed. Cir. 2012) ("If an applicant does not disclose structure for a means-plus-function term, the claim is indefinite.").

VLSI's proposed structure again is a generic "switching circuit." As discussed above, however, no figure or disclosed embodiment describes using a "switching circuit" for any

purpose—including to decouple the first power supply line from the second line of memory cells. DX-6 [Apsel] ¶¶ 91-92. Thus, VLSI's proposed construction should be rejected.

IV. U.S. PATENT NO. 7,606,983

A. Background

The '983 patent is directed to ordering processor requests to access shared memory. As the patent's background explains, digital electronic devices typically include multiple processors or "requestors" (in blue) all sharing access to common resources, such as main memory (in green):

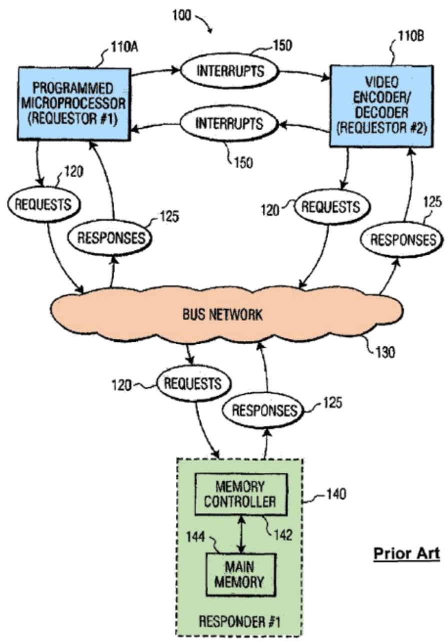

FIG. 1

DX-3 ['983 patent], 1:19-41, Fig. 1 (coloring added). Each requestor independently sends its own requests for access to the main memory over bus network 9 (in orange), and a memory controller or "responder" (in green) (1) resolves any contentions that occur when there are multiple active requests for access to the memory, (2) determines the order in which the multiple competing requests are to be performed, and (3) presents the requests to the memory one at a time. *Id.*

The patent further explains that the prior art typically used bus network protocols to implement policies for ordering "transactions" (i.e., an access request and its associated response).

Id., 2:19-30. These included, for example: (1) first-in-first-out ("FIFO") policies, in which a first transaction is completed before the next transaction starts; and (2) FIFO policies for transactions generated from the same requestor, with no ordering between transactions generated by different requestors. *Id.*, 2:34-54. The patent also notes that prior art systems could implement policies that imposed no constraints on when transactions are performed. *Id.*, 2:55-59.

The '983 patent purports to differ from the prior art by having the requestors (rather than the memory controller) determine transaction ordering. *Id.*, 4:57-5:11. To do so, the requestors include in each access request (1) an "indication" of whether the occurrence of the access request is to be performed in a sequential order; and (2) if so, a separate second "indication" of the specified order. *Id.*; DX-12 [Appeal Br.] at 5 (applicant asserting during prosecution that the patent's use of "two ***different*** indications" distinguished the claimed invention from prior art systems). For example, the patent notes that a requestor can provide these two indications to the responder as "signals" in the access request. DX-3 ['983 patent], 6:26-34.

B. Disputed Term: "an indication of a [the] specified order" (claims 1, 9, 11)

Intel's Construction	VLSI's Construction
"a second, different indication that indicates a | the specified order"	Plain meaning

Consistent with the applicant's characterization of the purported invention—both in the patent and during prosecution—the '983 claims require each access request to include two separate "indications": (1) "an indication of whether or not this occurrence of the access request [is to be performed in a sequential order/has a specified order] among other occurrences of the access request"; and (2) if so, "an indication of [a/the] specified order." DX-3 ['983 patent], 11:31-35, 12:16-20, 12:46-51. The parties dispute whether the second "indication" must be different from the first "indication," as Intel contends, or requires no construction, as VLSI argues. Intel's construction should be adopted because it is based on the plain language and gives effect to the

applicant's express representations made during prosecution.

First, the plain claim language supports Intel's construction. Each independent claim requires ***two separate and distinct*** indications in the access request, each of which must include different information: (1) the first "indication" must indicate whether to perform the specific occurrence of the access request in a sequential or specified order among other occurrences; and (2) if so, the second "indication" must then indicate the specified order. *Id.*; DX-7 [Hagersten] ¶ 43.

The claim language also makes clear that the second "indication" depends on the first "indication." For example, in claims 1 and 9, the access request includes the second "indication of a specified order" only "***if***" the first indication indicates that this occurrence of the access request is to be performed in sequential order. DX-3 ['983 patent], 11:34-35; *id.*, 12:19-20 ("***if so***, an indication of a specified order"). Similarly, in claim 11, the second "indication of the specified order" is "includ[ed]" "in those occurrence[s] of the access request ***that ar[e] ordered***." *Id.*, 12:49-50. This language—stating that the second "indication" is ***contingent on, and in addition to***, the first "indication"—further confirms that the two "indications" are separate and distinct from each other. *Id.*, 5:3-6, 6:26-32; DX-7 [Hagersten] ¶ 44.

Second, Intel's construction also reflects the applicant's express representations made to distinguish prior art. During prosecution, the Examiner repeatedly rejected claims as anticipated by a patent ("Weber") disclosing a single bit that the Examiner initially treated as satisfying both claimed "indications." After several unsuccessful attempts to distinguish Weber, the applicant filed an appeal brief arguing that the claims required "two ***different*** indications":

> Appellant's ***claimed invention*** includes ***two different indications***, each of which are provided by the access request. A first indication is related to whether the access request is to be performed in a sequential order among other occurrences of the access request. ***A second indication is used conditionally relative to the first***

> ***indication*** and is related to an indication of a specified order. The Weber reference does not teach that access requests contain these ***two different indications***.

DX-12 [Appeal Br.] at 5; *see also* DX-5 [Am. Appeal Br.] at 2 (identifying two different elements in the specification as the "indications"). The applicant further emphasized that the claimed invention required the second "indication" (specifying the order) to be ***in the access request*** itself:

> Weber teaches that for requests that are part of the same global group, the requests are serviced based upon the branch that was least recently serviced by the arbitration controller.… To accommodate this type of FIFO ordering, Weber's access requests specify the group to which they belong. This global group ***does not, however, specify the order*** in which they are to be performed ***nor does it indicate whether they are to be performed in a sequential order*** among other occurrences of the access request. ***The order in which Weber's access requests are serviced is left up to the arbitration controller rather than being identified in the access request itself***.

Id.; DX-7 [Hagersten] ¶¶ 45-46.

When the Examiner responded that Weber's single bit "teaches the two different indications," DX-14 [Reply Br.] at 6, the applicant argued that the claim language required "two ***different*** indications"—i.e., two separate indications that each convey different information:

> The Examiner's analysis is improper as it conveniently ignores the claim language supporting the second indication. The claim language immediately following the Examiner's selected expert [sic] is "and, if so, an indication of a specified order". ***Thus, the claim language explicitly recites a first indication of whether to perform in a sequential order and a second indication of the specified order. Accordingly, the Examiner's interpretation improperly ignores claim limitations directed toward two different indications***.

Id. at 4; DX-7 [Hagersten] ¶ 47. After receiving those representations, the Patent Office overruled the Examiner's rejections and issued the claims. DX-13 [Decision] at 7-8.

Accordingly, the applicant's statements confirm that the plain meaning of "an indication of a [the] specified order" requires two separate and distinct indications, just as Intel has proposed. *See Iridescent Networks, Inc. v. AT&T Mobility, LLC*, 933 F.3d 1345, 1352–53 (Fed. Cir. 2019) ("Any explanation, elaboration, or qualification presented by the inventor during patent

examination is relevant, for the role of claim construction is to 'capture the scope of the actual invention' that is disclosed, described, and patented.").[2]

VLSI asserts that the limitations reciting "an indication of a [the] specified order" need no construction. But at the very least, Intel's construction is needed to make clear to the jury that the plain meaning requires two different indications—as confirmed by the plain claim language and the applicants' representations made to secure issuance of the asserted claims.

V. U.S. PATENT NO. 7,793,025

A. Background

The '025 patent is directed to a technique for prioritizing requests—called "interrupts"—that a processor receives from other system components. DX-4 ['025 patent], 1:13-16 (describing how an "interrupt" is a request sent to a processor "from hardware indicating the need for attention" or "[from] software indicating the need for a change in execution"). The interrupt "causes the processor to save its state of execution" and "begin execution of an interrupt handler," which services the interrupt. *Id.*, 1:16-23. For example, a processor that is asleep might receive interrupts from an external device (e.g., from a keyboard when keys are pressed) that cause the processor to resume operation. *Id.*, 1:27-29; DX-8 [Farrens] ¶ 21.

Because a computer processor may receive multiple interrupts nearly simultaneously, and because some interrupts are more urgent than others, it was known in the prior art for conventional computer systems: (1) to assign priority levels to different pending interrupts; and (2) if multiple

[2] If VLSI incorrectly argues that "plain meaning" does not require two separate indications, the applicant's statements expressly and unmistakably disclaimed a single "indication" meeting both requirements. *See, e.g., Comput. Docking Station*, 519 F.3d at 1374 ("[A] patentee may limit the meaning of a claim term by making a clear and unmistakable disavowal of scope during prosecution…. A patentee could do so, for example, by clearly characterizing the invention in a way to try to overcome rejections based on prior art.").

interrupts were pending simultaneously, for a processor to handle them in order of their assigned priorities. *Id.*, 1:29-37; DX-20 [Practical Introduction to Elec. Circuits] at 467 ("Processors usually have more than one interrupt input, each having a different priority level."). As was known in the prior art, memory devices called "registers" can store pending interrupts, as well as their interrupt priority levels. DX-4 ['025 patent], 2:5-12; DX-19 [Computers as Components] at 121 ("[T]he CPU stores in an internal register the priority level of [the] interrupt."). It was also known that computers could operate with different modes or contexts, and that interrupt prioritization levels may change when the operating mode or context changes (e.g., when a system wakes from sleep mode or begins executing a different software program). DX-8 [Farrens] ¶ 24.

The '025 patent does not claim to invent interrupts or the idea to change from one set of interrupt priority levels to another when a system's operating mode or context changes—both of which the patent admits were known. DX-4 ['025 patent], 1:29-37, 1:63-2:13. Instead, the patent alleges that prior art techniques for switching between different sets of interrupt priority levels were too slow because the prioritization changes were "controlled by software." *Id.*, 2:5-11, 2:14-18; DX-8 [Farrens] ¶ 27.

The patent claims to improve over a software-based approach in just one way: by using an entirely ***hardware-based*** technique for switching among sets of interrupt priority levels upon a system mode or context change. DX-4 ['025 patent], 2:46-50 (explaining "the present invention" involves "a ***hardware*** mechanism … to adaptively prioritize interrupts based on the current mode or context of a corresponding processor"), Title ("***Hardware Managed*** Context Sensitive Interrupt Priority Level Control"); DX-8 [Farrens] ¶¶ 28-32. To do so, the patent describes a specific

arrangement of hardware components with two key features.[3]

First, the patent uses "interrupt priority registers," which are hardware components used to "stor[e] priority values." DX-4 ['025 patent], 2:64-67, Fig. 4.[4] The patent explains that an interrupt priority register stores a priority for ***every possible interrupt that might be received***. DX-4 ['025 patent], 3:3-7 ("***Each entry*** in each interrupt priority register corresponds to ***a possible pending interrupt request*** ...."); *id.*, 2:64-67 (similar); DX-8 [Farrens] ¶ 29.

Second, instead of using just a single interrupt priority register that is rewritten by software when the system mode or context changes (as in the prior art), the patent uses "multiple interrupt priority registers," one for "each context" or system mode. DX-4 ['025 patent], 3:1-3, 3:12-18 (describing "providing a separate interrupt priority register for each context"). To switch between these multiple hardware interrupt priority registers, the '025 patent uses a standard hardware "switching" or "multiplexer circuit[]." *Id.*, 7:58-61. This multiplexer circuit acts like a railroad switch; (1) in a first system mode, it couples a first interrupt priority register to the priority encoder (which provides interrupts to the processor), and (2) in a second system mode, it couples the second interrupt priority register to the priority encoder (and decouples the first interrupt priority register from the priority encoder). *Id.*, 9:47-54 ("[A] multiplexer circuit is provided … [that] selects and provides priority level information from one of the interrupt priority storage devices at an output in response to the mode control signal."); DX-8 [Farrens] ¶ 31.

In this manner, the '025 patent claims to improve over prior art systems that used software to rewrite a single interrupt priority register with new priorities upon a system mode or context

---

[3] Although the patent also discusses other peripheral features well known in the prior art, such as "masking" (a known technique for enabling and disabling interrupts), those features are not implicated by any claim term in dispute here.

[4] The '025 patent does not claim that using interrupt priority registers to store priority values for interrupts was new. *Id.*, 2:5-11 (describing use of interrupt priority registers in prior art).

change—by disclosing an ***entirely hardware-based*** system that switches among multiple distinct mode-specific ***hardware*** interrupt priority registers upon a system mode or context change. DX-4 ['025 patent], Abstract (describing the claimed invention as directed to "an interrupt circuit (300) that multiplexes (324) a plurality of interrupt priority registers (321, 322) based on the current context of the system"), Fig. 3; DX-8 [Farrens] ¶¶ 30-32.

B. Disputed Term: "storage device for storing priority level information" (claims 1, 9) / "sets of priority levels in ... storage devices" (claim 17)

Intel's Construction	VLSI's Construction
"hardware for storing priority level information that is not rewritten by software when the system changes mode or context" / "hardware for storing priority levels that are not rewritten by software when the system changes mode or context"	Plain and ordinary meaning

The parties dispute whether the claimed "storage devices" that store priority level information must consist of hardware that is not rewritten by software when the system changes mode or context, as Intel contends, or whether no construction is required, as VLSI maintains. As detailed below, Intel's proposed construction is required given the clear disclosures in the patent and prosection history that repeatedly criticize and claim to improve over software-based interrupt priority systems through a wholly hardware-based approach described as "the present invention."

First, the '025 patent makes clear that interrupt priority level information must be stored in ***hardware*** storage devices, and disclaims using software to rewrite this information when the system mode/context changes. Indeed, the patent itself is titled "***Hardware Managed*** Context Sensitive Interrupt Priority Level Control," and opens by explaining that "***the present invention*** relates to an interrupt controller that includes context sensitive interrupt prioritization ***controlled by hardware***." DX-4 ['025 patent], Title, 1:9-11; DX-8 [Farrens] ¶¶ 35-37. This description alone requires Intel's construction. *See Pacing Techs., LLC v. Garmin Int'l, Inc.*, 778 F.3d 1021, 1025 (Fed. Cir. 2015) ("When a patentee describes the features of the 'present invention' as a whole, he

alerts the reader that this description limits the scope of the invention."); *Verizon Servs. Corp. v. Vonage Holdings Corp.*, 503 F.3d 1295, 1308 (Fed. Cir. 2007) (same).

Second, the patent repeatedly disparages, and purports to improve upon, systems that use software to rewrite the interrupt priority levels stored in a register when the mode changes. DX-4 ['025 patent], 3:19-22; *id.*, 6:63-7:3 (describing "a latency penalty … associated with having frequent ***software*** updates of interrupt settings"). In fact, using software to rewrite priority levels is precisely what the patent claims to avoid:

> [I]nterrupt priority levels have ***conventionally*** been ***controlled by software***, which means that any changes in the prioritization of interrupts (which can occur when the system changes mode or context) requires additional time and programming complexity to switch the prioritization ***by re-writing the interrupt priority registers*** …. ***Accordingly, there is a need for an improved interrupt controller that allows interrupts to be quickly and efficiently prioritized***.

Id., 2:5-16; DX-8 [Farrens] ¶¶ 38-39.

The specification is saturated with similar reiterations that avoiding reliance on software to re-write priority level information is the keystone of the '025 patent:

- "By ***eliminating the requirement for software*** to manage the priority level changes before a context switch, interrupt processing can proceed with reduced latency and lower power consumption." DX-4 ['025 patent], 3:12-22.
- "With the disclosed arrangement, the interrupt circuit 300 provides a ***hardware-based prioritization of interrupt requests*** that may be adjusted to provide new prioritization when a context switch occurs. By having interrupt priority registers for each mode followed by a multiplex selection circuit that is ***controlled directly through hardware*** based on the system's context, a hardware mechanism is provided that tracks priority levels for each mode and automatically switches priorities upon a change to a new

mode. ***The disclosed priority switching mechanism*** allows interrupt priority levels for each mode to be switched ***without software intervention***, thus reducing the overall latency of changing context, and also reducing the workload of the processor." *Id.*, 8:48-60.

- "The interrupt priority register 212 may be implemented using any desired storage mechanism to store an interrupt value for each potential interrupt, including but not limited to a plurality of transparent latches, a multi-bit register, a read-only memory, hardwired gates, an application specific integrated circuit (ASIC), ***or any other fully hardware controlled circuit***." *Id.*, 6:3-9.
- "***[T]here is a latency penalty, as well as processor overhead,*** associated with having frequent ***software updates*** of the interrupt settings." *Id.*, 6:63-65.

The specification's repeated, consistent, and clear criticisms of storage devices that are rewritten by software, and its claims to improve over those systems via a "fully hardware controlled circuit," separately require Intel's construction. *See Chicago Bd. Options Exch., Inc. v. Int'l Sec. Exch., LLC*, 677 F.3d 1361, 1372 (Fed. Cir. 2012) ("[T]he specification goes well beyond expressing the patentee's preference …, and its repeated derogatory statements about [features of the prior art] reasonably may be viewed as a disavowal …."); *Openwave Sys., Inc. v. Apple Inc.*, 808 F.3d 509, 513 (Fed. Cir. 2015) ("[I]t is difficult to envisage how, in light of the repeated disparagement of mobile devices with 'computer modules' discussed above, one could read the claims of the patents-in-suit to cover such devices.").

Third, consistent with its criticisms of software-based systems, the '025 patent claims to improve over those systems in only one way—by repeatedly teaching that the claimed interrupt priority levels are ***entirely implemented in, and controlled by, hardware***. DX-4 ['025 patent],

6:3-9 (describing how "[t]he interrupt priority register" is implemented in a "fully hardware controlled circuit"); *id.*, 1:9-11 ("[T]he present invention relates to an interrupt controller that includes context sensitive interrupt prioritization ***controlled by hardware***."); *id.*, 2:48-49 ("[A] ***hardware mechanism*** is provided to adaptively prioritize interrupts …."); *id.*, 3:45-48 ("[D]ata processing system 10 includes … ***one or more hardware devices*** …."); *id.*, 8:48-51 ("[I]nterrupt circuit 300 provides a ***hardware-based*** prioritization of interrupt requests …."); *id.*, 9:24-34, 9:42-47, Fig. 4 (steps 403, 404) (using terms "interrupt priority storage devices" and "***hardware*** registers" interchangeably); DX-8 [Farrens] ¶¶38-39 .

By contrast, no embodiment discloses using software—consistent with the patent's criticisms of a software-based approach. This too confirms Intel's construction. *See Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 653 F.3d 1296, 1305 (Fed. Cir. 2011) (limiting claims to embodiments with a particular feature where the specification "only disclose[d] embodiments [with that feature]" and "expressly distinguish[ed] the invention from the prior art based on [that] feature"); *SciMed Life Sys., Inc. v. Advanced Cardiovascular Sys., Inc.*, 242 F.3d 1337, 1341 (Fed. Cir. 2001) ("Where the specification makes clear that the invention does not include a particular feature, that feature is deemed to be outside the reach of the claims of the patent, even though the language of the claims, read without reference to the specification, might be considered broad enough to encompass the feature in question.").

Finally, beyond these many statements in the specification, the file history also supports Intel's construction by unequivocally disavowing a software-based solution and representing to the Patent Office that the claimed storage devices are (1) entirely hardware-based and (2) not rewritten by software when the system changes mode or context.

Specifically, during prosecution, the Patent Office initially rejected claims 1, 9, and 17 as

unpatentable over a prior art reference ("Kershaw"). DX-15 [Office Action] at 2-3. The applicant overcame these rejections by arguing there is a "very real difference" between the alleged '025 invention and Kershaw because the latter "uses the *operating system* ***software*** to set the interrupt priority values," whereas the '025 approach requires that "interrupt priority level information for each requested interrupt request is stored in '*interrupt priority storage* ***devices***.'" DX-16 [Response] at 9 (italics in original). The applicant further argued that (1) "Kershaw merely describes ***the conventional software-controlled interrupt handling approaches described and distinguished by Applicants in the 'Background of the Invention' section***," and (2) the '025 patent "***eliminat[es] the requirement for software management*** of priority level changes during context switches." *Id.* at 7, 9 (citing'025 application excerpt criticizing prior art where "interrupt priority levels have ***conventionally been controlled by software***"); DX-8 [Farrens] ¶¶ 40-42.

This argument swayed the Patent Office, which allowed the claims in "further … view of Applicant's persuasive argument that 'Kershaw uses the *operating system software* to set the interrupt priority levels, whereas claim 1 specifies that different interrupt priority level information for each request is stored in '*interrupt priority storage devices*,' depending on the system mode.'" DX-17 [Allowance] at 2 (italics in original).

Thus, the applicants obtained the '025 claims by representing that the claims are limited to hardware-based systems. VLSI is bound by those representations here, and cannot seek to cover systems that use software to rewrite priority level information. *See Comput. Docking Station Corp.*, 519 F.3d at 1376; *Krippelz*, 667 F.3d at 1266; *Aylus Networks, Inc. v. Apple, Inc.*, 856 F.3d 1353, 1359 (Fed. Cir. 2017) (parties cannot "recaptur[e] through claim interpretation [the] specific meanings" disavowed during prosecution); *CVI/Beta Ventures, Inc. v. Tura LP*, 112 F.3d 1146, 1158 (Fed. Cir. 1997) ("[T]hrough statements made during prosecution … an applicant … may

commit to a particular meaning for a patent term, which meaning is then binding in litigation.”).

C. Disputed Term: “priority level information associated with a [first/second] system mode for each of the one or more interrupt requests” (claim 1) / “priority level information associated with a [first/second] system mode” (claim 9)

Intel’s Construction	VLSI’s Construction
“priority level information associated with a [first/second] system mode for each of the one or more potential interrupt requests”	Plain and ordinary meaning

The parties dispute whether these limitations should be construed to make clear that the claimed “priority level information associated with a [first/second] system mode” is “for each of the one or more ***potential*** interrupt requests,” as Intel has proposed, or whether no construction is required, as VLSI contends. Intel’s construction captures the clear representations made—both in the ’025 patent and separately to the Patent Office during prosecution—that the claimed “mode-specific priority level information” must include priority level information for ***each potential*** interrupt request (rather than just for the pending interrupt requests that ***actually*** occur).

First, the patent repeatedly teaches that—unlike the “pending registers” that store “pending interrupt signals”—“interrupt priority registers” store “sets of priority levels” that include priority information for ***each potential interrupt request*** that ***might*** be received. *E.g.*, DX-4 [’025 patent], Figs. 2-4, 2:63-67, 5:65-6:18, 7:53-57, 8:60-9:10, 10:6-12, 10:51-55; DX-8 [Farrens] ¶¶ 44-47. For example, in introducing its purportedly novel way of storing priority level information, the patent explains that “[o]ne or more interrupt priority registers are also provided for storing priority values ***corresponding to each of the potential pending interrupt requests***, thereby enabling prioritization of the enabled, pending interrupt requests.” DX-4 [’025 patent], 2:63-67. Likewise, the specification further states that “***[e]ach entry*** in each interrupt priority register ***corresponds to a possible pending interrupt request***”—not a request that has already been generated by or received from an interrupt source. *Id.*, 3:3-5. Similar explicit statements that priority level

information must be associated with ***potential*** interrupt requests pervade the specification. DX-8 [Farrens] ¶¶ 44-47.[5]

Accordingly, because the ’025 patent "‘repeatedly and consistently’ characterizes [the] claim term in a particular way, it is proper to construe the claim term in accordance with that characterization." *GPNE Corp. v. Apple Inc.*, 830 F.3d 1365, 1370 (Fed. Cir. 2016).

Second, during prosecution, to overcome the Kershaw prior art reference discussed above, the applicant expressly represented that ***all claims*** (including claims 1 and 9) require storing priority level information for ***each potential*** pending interrupt request—not merely for actual pending interrupt requests. Specifically, the applicant argued that the claimed invention was distinguishable because Kershaw allegedly failed to disclose that the interrupt priority registers stored "***each potential*** pending interrupt request."[6] DX-16 [Response] at 7. In support of that distinction, the applicant referenced claim 17's express language requiring "a priority level for <u>each potential pending interrupt request</u>.’" *Id.* In contrast, the applicant argued that each of the "priority registers" in Kershaw only stored "<u>some</u>" potential interrupts. *Id.* at 9-10.

Importantly—as the Examiner recognized—even though claims 1 and 9 do not contain the word "potential" (unlike claim 17), the applicant made clear that ***every claim*** in the patent requires

[5] *Id.*, 10:51-55 (explaining how "each set of priority levels … specifies a priority level for each ***potential*** pending interrupt request"); *id.*, 5:65-67 ("[I]nterrupt priority register 212 may store predetermined priority values for each of the ***potential*** interrupts …."); *id.*, 6:3-6 ("The interrupt priority register 212 may be implemented using any desired storage mechanism to store an interrupt value for each ***potential*** interrupt ….").

[6] Kershaw describes a technique for mapping interrupts to different priorities depending on whether the system is operating in a secure mode or in a non-secure mode. DX-26 [Kershaw], Abstract. Like the ’025 patent, Kershaw discloses storing interrupt priority values in priority registers and using different sets of priority values for different modes. *Id.* at 2:3-24. However, according the applicant, instead of pre-storing priority levels for every potential interrupt request, Kershaw only stored some priority levels and used software to rewrite them as needed. DX-16 [Response] at 9-10.

storing "***each of the potential*** pending interrupt requests." *Id.* at 7 (explaining for claims 1 and 17 that "interrupt priority registers are provided for storing priority values ***corresponding to each of the potential pending interrupt requests***"); *id.* at 11 (explaining for claim 9 that, "[a]s described in the Application, the interrupt priority registers are provided for storing priority values ***corresponding to <u>each</u> of the potential pending interrupt requests***"); *id.* at 7-8 (emphasizing for ***all claims*** the patent's disclosure that "[o]ne or more interrupt priority registers are also provided <u>for storing priority values ***corresponding to each of the potential pending interrupt requests***</u>," and arguing that ***all claims*** differed from Kershaw based on the specification's explanation that "[e]ach entry in each interrupt priority register corresponds to a ***possible*** pending interrupt request" (underlining in original)); DX-8 [Farrens] ¶¶ 48-52.

Thus, when arguing to the Patent Office why claims 1 and 9 should be granted, the applicant represented that those claims require the interrupt priority registers to store "priority values corresponding to each of the ***potential*** pending interrupt requests." DX-4 ['025 patent], 2:65-66. And the Patent Office expressly relied on this argument in deciding to issue the claims. DX-17 [Allowance] at 2 (citing applicant's "persuasive" point that "[a]s described in the Application, the interrupt priority registers are provided for storing priority values corresponding to each of the ***potential*** pending interrupt requests").

In light of this clear intrinsic evidence, VLSI cannot abandon the meaning of the disputed limitations that the applicant represented to the Patent Office to obtain the claims. To prevent that unfair and improper result, Intel's construction should be adopted. *See Aylus Networks*, 856 F.3d at 1359 ("[T]he doctrine of prosecution disclaimer ensures that claims are not construed one way in order to obtain their allowance and in a different way against accused infringers.").

D. Disputed Term: "providing a plurality of interrupt priority storage devices … and providing a plurality of interrupt priority storage devices …"

(claim 1)

Intel's Construction	VLSI's Construction
Indefinite	Definite

The parties dispute whether claim 1—which contains an incoherent repetition of elements—should be deemed indefinite, as Intel contends, or is "definite," as VLSI has urged (without specifying any actual proposed meaning). Because the limitation renders the claim nonsensical in several respects, Intel's position should be adopted.

Specifically, claim 1 recites both of the following "providing" limitations, with the only differences shown as underlined:

First "providing" limitation	Second "providing" limitation
providing a plurality of interrupt priority storage devices comprising a first interrupt priority storage device for storing priority level information associated with a first system mode, and a second interrupt priority storage device for storing priority level information associated with a second system mode; and	providing a plurality of interrupt priority storage devices comprising a first interrupt priority storage device for storing priority level information associated with a first system mode <u>for each of the one or more interrupt requests</u>, and a second interrupt priority storage device for storing priority level information associated with a second system mode <u>for each of the one or more interrupt requests</u>; and

Both limitations require (1) "***a plurality*** of interrupt priority storage devices," (2) a ***first*** interrupt priority storage device for storing priority level information associated with a first system mode," and (3) "a ***second*** interrupt priority storage device for storing priority level information associated with a second system mode." It is unclear, therefore, whether the claim requires ***two*** priority storage devices (if the "first" and "second" devices are the same), or ***four*** devices altogether (two separate "pluralities"). DX-8 [Farrens] ¶¶ 53-55.

Moreover, assuming the claim requires two separate pluralities of interrupt storage devices (four devices altogether), as it is written, the claim has insurmountable antecedent basis problems. The last limitation of the claim imposes requirements for "***the*** plurality of interrupt priority storage

devices"; yet, it is impossible to determine which "plurality of interrupt priority storage devices" serves as the antecedent basis for this limitation. DX-8 [Farrens] ¶¶ 56-57.

Because this illogical redundancy makes no sense on several levels, the limitation should be construed as indefinite. *See Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898, 899, 910 (2014) (holding limitations are indefinite if they fail to "inform those skilled in the art about the scope of the invention with reasonable certainty," because that lack of disclosure fails to "appris[e] the public of what is still open to them"); *Trustees of Columbia Univ. in City of N.Y. v. Symantec Corp.*, 811 F.3d 1359, 1367 (Fed. Cir. 2016) ("The claims are nonsensical in the way a claim to extracting orange juice from apples would be, and are thus indefinite.").[7]

VI. U.S. PATENT NO. 7,523,373

A. Background

The '373 patent addresses aspects of the well-known challenge in integrated circuit design of balancing "tradeoffs between performance and power." DX-5 ['373 patent], 1:12-18. As was well understood when the '373 patent was filed, it is ideal for processors to operate "at maximum voltage and frequency when peak performance is required," and at other times to operate "at low voltage and frequency to reduce power consumption." *Id.* But because "different types of circuitry within a data processing system may have different ranges of allowable operating voltages," *id*, 1:23-25, "the processor may be able to operate at a lower voltage than is possible for the memory," *id.*, 2:5-7; DX-9 [Sylvester Decl.] ¶ 26.

[7] VLSI has not proposed a construction for this term that would fix the claim's incoherency, and the claim cannot be re-written to solve its indefiniteness, even if VLSI claims that a drafting error occurred. *See Lucent Techs., Inc. v. Gateway, Inc.*, 525 F.3d 1200, 1215 (Fed. Cir. 2008) ("[C]ourts may not redraft claims to cure a drafting error made by the patentee, whether to make them operable or to sustain their validity.").

To enable a processor to operate at voltages below the memory's minimum operating voltage, the patent discloses a specific way to provide an alternative voltage source for the memory: (1) when a processor is operating at a voltage ***above*** the memory's minimum operating voltage, a ***single*** regulated voltage is provided to ***both*** the processor logic and the memory; (2) however, if that voltage falls ***below*** the memory's minimum operating voltage, then the ***alternative*** (higher) regulated voltage is provided to the memory, while the original (lower) voltage is still provided to the processor logic. *E.g., id.*, Abstract, 3:30-45, claims 1, 16.

B. Disputed Term: "means for providing the operating voltage to the memory at a value at least as great as the minimum operating voltage in response to the operating value selected by the processor being below the minimum operating voltage" (claim 14)

Intel's Construction	VLSI's Construction
Indefinite	**Function**: "providing the operating voltage to the memory at a value at least as great as the minimum operating voltage in response to the operating value selected by the processor being below the minimum operating voltage" **Structure**: "power supply selector, charge pump, scalable voltage regulator, or equivalents thereof"

The parties dispute whether the "means for providing the operating voltage to the memory …" limitation of claim 14 (which depends from claim 9) is indefinite because it recites a nonsensical function (requiring two separate "operating voltages" to the same memory) and no corresponding structure for performing that illogical function, as Intel contends, or whether it should be construed to have the claimed function and incomplete structure that VLSI proposes. Intel's proposed construction should be adopted for the reasons below.

Function: As an initial matter, the claimed function itself is indefinite. On its face, the language of claim 14 requires a "means for ***providing the operating voltage to the memory***." But claim 9 (from which claim 14 depends) claims an "integrated circuit" that already includes a separate component—a "power supply selector"—that "supplies … ***the operating voltage of the***

memory." DX-5 ['373 patent], claims 9, 14. These conflicting requirements are summarized in the table below:

Claim	Provided to "the memory"	Component Providing the Operating Voltage	When the Operating Voltage is Provided to the Memory
9	"the operating voltage"	"power supply selector"	"when the first regulated voltage is below the minimum operating voltage"
14	"the operating voltage"	"means for providing the operating voltage to the memory"	"in response to the operating value selected by the processor being below the minimum operating voltage"

Thus, dependent claim 14 requires ***two different*** sources to supply "***the*** operating voltage" to "***the*** memory." But it would be paradoxical to have two ***different*** components supply the ***same*** operating voltage to the ***same*** memory, especially under similar and non-exclusive conditions. DX-9 [Sylvester Decl.] ¶¶ 31-33. Accordingly, because the claim specifies a nonsensical function, the limitation is indefinite. *See Trustees of Columbia Univ.*, 811 F.3d at 1367.

Structure: This limitation is also indefinite because the patent fails to disclose a structure that performs the claimed function—which is not surprising given that the structure itself is illogical. As noted above, claim 14 requires the memory itself to include the "means for providing ***the*** operating voltage to ***the*** memory" ("the memory of claim 9 further comprising … means for providing …"), even though the "power supply selector" of claim 9 already must provide "***the*** operating voltage" to "***the*** memory." There is simply no structure that the '373 specification describes in which ***two different*** components provide the ***same*** operating voltage to the ***same*** memory. DX-9 [Sylvester Decl.] ¶¶ 34-37. This lack of disclosed structure also renders the limitation indefinite. *See Williamson*, 792 F.3d at 1352.

For its part, VLSI argues that the patent discloses a "power supply selector, charge pump, scalable voltage regulator, or equivalents thereof" as the corresponding structure. But that argument fails for multiple reasons.

First, as noted above, claim 9 requires the claimed "power supply selector" to supply "the operating voltage" to "the memory," while claim 14 expressly requires the "means for providing the operating voltage to the memory …" to be internal to the same memory. But the patent does not disclose any emobodiment in which a "power supply selector" supplies an operating voltage to a memory, and a separate structure in the memory also "provid[es] ***the*** operating voltage to the memory." DX-9 [Sylvester Decl.] ¶¶ 34-37. Nor would it make sense for a structure ***within*** the memory to provide something "***to*** the memory." *Id.* ¶ 38.

Second, the "power supply selector" described in the '373 specification does not perform the function that VLSI identifies for claim 14—either by itself or in combination with the "charge pump" and/or "scalable voltage regulator" that VLSI includes as the corresponding structure. Nowhere in the specification does a "power supply selector" provide the operating voltage to the memory at a particular value "in response to the operating value ***selected by the processor*** being below the minimum operating voltage." For each disclosed embodiment that mentions the "power supply selector," there is no discussion of (1) the "processor" selecting an operating value for the memory's voltage or (2) how other components act based upon such a selection.[8]

For example, in the first disclosed embodiment, which discusses power supply selector 21 of Figure 1, the operating voltage is determined by controller 28—not processor 16—and power supply selector 21 merely selects between two voltages based on a signal received from controller

[8] Some disclosed embodiments explain how the circuitry should respond to ***other*** actions taken by the processor, such as when the processor selects its operating ***state*** (which controller 28 uses to determine the frequency and operating voltage of the processor, and then uses that information to decide which regulated voltage the memory should receive). DX-5 ['373 patent], 8:16-47. But this embodiment appears to relate to claim 6, which requires that a "***controller*** … select[] an operating value for the operating voltage of the memory." Nothing in this text (or anywhere else in the patent) discloses structure clearly linked to either the "***processor*** … select[ing] an operating value ***for the operating voltage of the memory***," as claim 14 provides, or "respon[ding] to the operating value ***selected by the processor***," as this disputed limitation requires.

28. DX-5 ['373 patent], 2:55-57 ("Power supply selector 21 selects one of VDDmem and VDDlogic based on information provided by controller 28 via, for example, the memory control signals."); DX-9 [Sylvester Decl.] ¶ 40. Power supply selector 21 does not (1) receive an operating value for the operating voltage of the memory, (2) from ***the processor***, and (3) determine whether that operating value is lower than the memory's minimum operating voltage. DX-9 [Sylvester Decl.] ¶ 40. Indeed, the only time that the specification mentions an operating value selected by the processor for the memory, it merely parrots the functional language of this limitation—without identifying any structure to perform that full function. DX-5 ['373 patent], 11:24-30.[9]

Finally, as to the "scalable voltage regulator" and "charge pump" that VLSI identifies as part of the corresponding structure, neither component is linked to the claimed function that VLSI alleges. While the specification discloses multiple voltage regulators, and notes that one might be scalable, the voltage regulators merely provide a voltage dictated by other circuitry; they are not described as capable, by themselves, of "respon[ding] to the operating value selected by the processor being below the minimum operating voltage." Any comparison between an operating value and a minimum operating voltage is performed by other components. *E.g.*, *Id.*, 5:42-53; DX-9 [Sylvester Decl.] ¶ 43.

Moreover, the "charge pump" is only mentioned once in the specification, and identified there only as a means of boosting voltage as dictated by other circuitry—not as a means of determining whether some selected operating value is lower than the minimum operating voltage. DX-5 ['373 patent], 5:58-61; DX-9 [Sylvester Decl.] ¶ 44.

Thus, even if VLSI's claimed function is adopted, the limitation still would be indefinite

[9] The patent also does not disclose a "memory … comprising … ***a processor***," as claim 14 requires; the disclosed embodiments only discuss a processor ***external*** to the memory. *E.g.*, *id.*, Fig. 1.

because no disclosed structure is clearly linked to performing that function. *See Default Proof Credit Card Sys., Inc. v. Home Depot U.S.A., Inc.*, 412 F.3d 1291, 1299 (Fed. Cir. 2005) ("To meet the definiteness requirement, structure disclosed in the specification must be clearly linked to and capable of performing the function claimed by the means-plus-function limitation.").

Dated: October 30, 2019

Respectfully submitted,

/s/ J. Stephen Ravel
J. Stephen Ravel
Texas State Bar No. 16584975
KELLY HART & HALLMAN LLP
303 Colorado, Suite 2000
Austin, Texas 78701
Tel: (512) 495-6429
Email: steve.ravel@kellyhart.com

James E. Wren
Texas State Bar No. 22018200
1 Bear Place, Unit 97288
Waco, Texas 76798
Tel: (254) 710-7670
Email: james.wren@baylor.edu

Sven Stricker
Texas State Bar No. 24110418
KELLY HART & HALLMAN LLP
303 Colorado, Suite 2000
Austin, Texas 78701
Tel: (512) 495-6464
Email: sven.stricker@kellyhart.com

Attorneys for Intel Corporation

OF COUNSEL:

William F. Lee (*Pro Hac Vice)*
Louis W. Tompros (*Pro Hac Vice)*
Kate Saxton (*Pro Hac Vice*)
WILMER CUTLER PICKERING HALE
& DORR LLP
60 State Street
Boston, Massachusetts 02109
Tel: (617) 526-6000
Email: william.lee@wilmerhale.com
Email: louis.tompros@wilmerhale.com
Email: kate.saxton@wilmerhale.com

Gregory H. Lantier (*Pro Hac Vice*)
Amanda L. Major (*Pro Hac Vice*)
WILMER CUTLER PICKERING HALE
& DORR LLP
1875 Pennsylvania Avenue
Washington DC 20006
Tel: (202) 663-6000
Email: gregory.lantier@wilmerhale.com
Email: amanda.major@wilmerhale.com

CERTIFICATE OF SERVICE

I hereby certify that all counsel of record who are deemed to have consented to electronic service are being served with a copy of the foregoing document via the Court's CM/ECF system per Local Civil Rule CV-5(b)(1) on October 30, 2019.

/s/ J. Stephen Ravel
J. Stephen Ravel